



A threat can come
from anyone,
anywhere,
anytime.

Let us be your early warning system.

In a time when social media is the primary form of communication, the Social Sentinel® service helps you provide a safe environment for students, faculty, staff, and visitors—all while respecting the rights of everyone in your community.



We help you...

Identify social media threats from across the street and country.

Every alert stems from the Sentinel SearchSM Library. We apply over 450,000 behavioral threat indicators to over 1 billion public social media posts per day to surface threats on your behalf—regardless of who posts the threat and where he/she is located.

Maximize efficiency and minimize FTE engagement.

Social Sentinel works in the background reviewing billions of posts per week, alerting you to a small number of concerning posts associated with your community to assess and address.

Target threats, not people.

You can control how the system is being used by ensuring that rules are applied to identify threats. Show your community that you respect its rights to privacy and association by incorporating constraints that prevent users from targeting individuals or groups.

Mitigate risks to fulfill your Duty of Care.

Whether it's your parent community, board of directors, or insurance company, you can say, "We didn't wait to hear about the threat. We put a system in place to do everything we could to be notified of threats immediately, and had the details to begin assessment in near real time."

For a **demonstration** or **additional details** about our service, contact us at
800.628.0158 or **sales@socialsentinel.com**

Social Sentinel FAQ

Why does Social Sentinel exist?

Safety is a human right, and our need to be connected is a human drive. The explosive proliferation of social media finds some users posting their harmful intentions—publicly—towards the people and places where we learn, work and play. Our service provides clients with insights into those public threats to help maintain the safety of their communities. We do this with a commitment to protect our collective (and constitutional) rights to: privacy, freedom of association, and freedom of speech; a commitment found in the DNA of the design and delivery of our service.

What is Social Sentinel?

Since 2015, the Social Sentinel[®] service (Social Sentinel) has provided near real-time awareness of possible safety/security threats made on public social media. We built our service to assess public social media across multiple platforms to detect potential acts of harm or threats relevant to our clients.

Such threats could range from mass shootings to destruction of your property. Social Sentinel's service is an important safety and security tool among many needed today to create awareness of possible tragedies, and to provide a window of time to assess and respond accordingly. We believe that early intervention is a critical step in preventing such tragedies from becoming a reality.

Who uses Social Sentinel?

Any organization that creates a gathering of people with the intent to learn, work, or have fun. For example: schools, colleges, universities, sports teams, arenas, festivals, convention centers, corporate campuses, theaters, and more.

How does Social Sentinel find threats?

We assess public posts for matches against our proprietary Sentinel SearchSM Library to determine if they contain the language of harm. This pre-populated collection of more than a half million behavioral threat indicators was developed by experts in mental health, public safety, security, linguistics and data science. The Library continually evolves and adapts to address changes in language and expression under the guidance of our Librarian, linguists and data science team.



How does Social Sentinel work?

Our system has four stages:

Access: Social Sentinel has authorized access to over one billion public social media posts daily.

Process: We scan for more than a half million behavioral threat indicators found in our Sentinel SearchSM library.

Match: We map alerts to your team via our proprietary matching processes powered by unique characteristics of your assets.

Alert: Get alerts in near real time and process them through your existing threat assessment protocols.

How are alerts shared with a client?

Alerts can be sent via text and/or email to our clients' users, and they can access alerts through our online application. The recipients and their preferred delivery methods can be updated at any time.

Does Social Sentinel see private posts?

No, we do not identify alerts from private posts. We have developed our industry-leading standard to protect constitutional rights while allowing for meaningful, preventive public safety action.

Is Social Sentinel a monitoring, surveillance or investigation tool?

No, and this is an important distinction. Aligning with our commitment to protecting constitutional rights, we developed Social Sentinel to be a threat alert service that provides information about imminent safety and security issues. Our solution is not built as an investigative tool, and cannot be used as such.

Does Social Sentinel violate an author's right to privacy?

No. We believe that privacy is a human right as much as a constitutional right. Our service accesses only public social media, which can be seen by anyone, anywhere, anytime. If the author's account is marked as private, or they use a closed social network, our service does not have access.

Access to data from social media partners: Social Sentinel accesses data from our social media partners through specific, approved use cases and via permissions outlined in API Terms of Use. Social Sentinel does not use any of this data, including from Facebook and Instagram, for monitoring. For Instagram, use of the data is limited to public content that explicitly includes the client's named assets.

© 2018 Social Sentinel, Inc. All rights reserved. The information in this document is for informational purposes only and is not for the purpose of providing legal advice.

Social Sentinel

White paper | February 2018

7 Essential Qualities of Effective Social Media Threat Alert Services

7 Essential Qualities of Effective Social Media Threat Alert Services

Right now, critical information that could enhance your ability to identify risks, assess threats, and manage events is being shared over social media. Choosing a social media threat alert service for your district – regardless of its size – may seem daunting at first. To help you evaluate a provider thoroughly, keep the following seven essentials in mind during your selection process.

1. Deliver actionable alerts

Officials must be notified in real-time of an actionable alert. Resources are limited and district personnel must evaluate each alert as it comes in. To make this viable, the service that the administration selects must be adept at surfacing alerts that are important to district's, while minimizing the inclusion of content that is not of imminent importance.

When evaluating vendors, look for a provider that can demonstrate proficiency in the issues that impact district's and that has experience differentiating among content associated with an imminent issue, content that provides situational awareness, and content that is just "noise."

2. Avoid targeting and "fishing expeditions"

District officials must ensure that the service provider selected has created a method of discovering alerts in a manner that does not allow district employees to target an individual or group or to go on "fishing expeditions."

In October 2016 the ACLU published a series of opinion pieces¹ following an investigation into social media monitoring services being used by law enforcement agencies for surveillance purposes. This action resulted in the largest social media companies denying data access to some social media monitoring companies. As a result, some are no longer in business. The primary challenges that the ACLU raised were two-fold. First, the social media monitoring objections were allowing their end users to profile and/or surveil social media users. Second, they were allowing end users to enter whatever search parameter they wanted with no oversight or accountability. Some of the methods used included:

- Drawing a geofence around an area and watching all of the comments that came through (e.g., going on a fishing expedition);
- Entering into the social media monitoring company's keyword tool a specific phrase (in the ACLU's example "#blacklivesmatter) with no oversight or explicit reason; and
- Allowing end users to identify "influencers;" and then linking those "influencers" to other social media users within the social media platform and across other social media platforms.

¹ For more information see this blog post: <https://www.aclunc.org/blog/facebook-instagram-and-twitter-provided-data-access-surveillance-product-marketed-target>

By contracting with a threat alert company versus a social media monitoring company, you more likely are weeding out the service providers that facilitate this kind of behavior, but you need to make sure that the service is, in fact, created in a manner that does not facilitate the above behaviors.

To ensure that district personnel are unable to go on fishing expeditions and/or target individuals or groups, look for a service provider with a systemic mechanism to curb or eliminate this possibility. This should include, at a minimum:

- A comprehensive library of school-related public safety triggers. Less sophisticated service providers rely on keywords (which will produce a high volume of false positives). More sophisticated providers will be able to describe the layers of data science used. The goal is to have your team alerted to an issue rather than requiring them to surface an issue;
- The absence of the ability to draw a geofence and monitor the conversation inside;
- The absence of the ability to enter a keyword and monitor what people are saying around that keyword;
- A means to relate content to your organization so that alerts are narrowly focused; and
- A means to oversee what district personnel are doing within the service.

3. Mapping threats vs sifting through geofenced content

District officials need their service provider to be able to associate a threat with their community and the assets associated with that community. This is not an easy task.

In the past, social media monitoring companies relied upon a geofence to link content to an organization (typically a brand, service, or product). It's easy, cheap, and provides a reasonable sample of data available. However, a district that wants to be alerted to a threat in near real-time will not be well-served by this process because:

- Less than 5% of all publicly available social media content has a geocode attached to it and a geocode must be attached for the content to show up within the geofence. School officials looking for a needle in the haystack are unlikely to find it if only 5% of the content produced from within that geofence is being included in the service's analysis process; and
- A threat can come from anyone who may be posting from anywhere. It is unlikely that the threat will be posted by someone currently residing within the selected geofence.

Successful service providers must be able to demonstrate best efforts to deliver school-specific alerts without relying on a geofence.

4. The importance of employee oversight

District officials have a duty to oversee what their employees are and are not doing with the services made available to them to perform their administrative tasks. Further, should an unforeseen incident occur, it is not unusual for officials to be required to provide insurance companies, trustees, and judicial officials with documentation showing how employees did and did not use such services throughout that incident.

Service providers must be able to demonstrate how district administration can oversee employees on an ongoing and ad hoc/as needed basis.

5. Publicly available content vs closed networks and apps

All content included in the service must be publicly available. While district's may have a service that monitors content (email, web searches, etc.) that a student or staff member has produced using school provided devices, when receiving alerts to threats shared on social media outside of those devices and accounts, that content must not come from private environments such as closed networks, sites, or messaging apps.

6. Providers need licenses to access content

Officials must ensure that the chosen threat alert company has the explicit right to access the social media companies' content for safety and security purposes.

Social media companies provide a variety of different content access licenses to businesses including threat alert and social media monitoring companies. The most common license allows brands to gauge how people feel/talk about their company, products, and/or services. Social media companies are selective in who they work with for safety and security access as there is the potential that the end user/recipient of the alert (e.g., a risk mitigation official, a police officer) may take action against the content poster/creator (the social media company's customer) in a manner that may impact their customer's constitutionally protected rights. As discussed earlier, the ACLU tracks this issue closely.

The few threat alert companies and the more traditional social media monitoring companies that do have this explicit license must pass rigid use-case review

processes and submit to ongoing vetting. There are a number of social media companies that currently operate without this license. They use licenses approved for other purposes or they "scrape" data to control costs and subvert this license altogether. In the process, they are violating social media companies' Terms of Use and subjecting their company to loss of access to data, and, in some cases, criminal and/or civil action.

Require respondents to affirm that they are operating with licenses granting them the ability to access social media data for safety and security purposes.

7. Providers need licenses to distribute content to public safety teams respectively

District officials must ensure that the threat alert company chosen has the explicit license to pass content that triggers a safety and/or security alert to a district or district related public safety team member. The license that a threat alert company or social media monitoring company receives is a two-step process. The first grants access to the type of content associated with a safety/security issue (discussed above). The second grants the right to pass the content to certain types of government and/or safety officials. Some threat alert and social media monitoring companies, for instance, are able to secure a license for passage to a corporate security team but not to a district or district related safety official based on how their service works. There is an additional layer of scrutiny associated with passage to end users who work for various safety and security entities.

When evaluating vendors, verify that the threat alert or social media monitoring company responding has, in fact, secured the explicit license to pass content to all of the members of your organization who will be receiving alerts (e.g., district safety/security team members, SRO, risk mitigation officials) for safety and security content.